Numerical Linear Algebra and Optimisation

Credits: 7.5hp

Time: Beginning in Period 2, Fall 2024 and then regularly every second year.

Course structure: A series of lectures, which may be pre-recorded or live, in combination with seminars where the material is discussed.

Examination: The course will be examined through assignments and project work, including oral presentations to the group.

Level: The course is targeted to graduate students with some background in mathematics and scientific computing.

Content: The course is intended to cover prominent topics in numerical linear algebra and optimisation. Specifically the following areas and related topics will be included.

Numerical Linear Algebra

- Basic matrix theory: various types of matrices (e.g., symmetric/Hermitian, unitary, normal, positive definite, indefinite, reducible/irreducible, etc.). We will also cover topics including Schur and spectral decomposition, Jordan canonical form, LU, LLT, block LU/LDU, etc.
- Representation of dense and sparse matrices, and matrix libraries
  - Regular column- or row-wise storage
  - Compressed sparse row, quadtree, etc
  - Matrix-free/on-the-fly computed representations
  - BLAS and LAPACK
- Krylov subspace methods for eigenvalues and linear systems
  - Arnoldi
  - Lanczos
  - Conjugate gradient
  - GMRES
  - Methods for Least Squares problems
- Stability and backward error analysis of some methods
- Functions of matrices: f(A)

Optimisation

- Introduction: unconstrained vs constrained, global vs local, derivative-free v/s derivative-based, first order v/s second order
- Convex optimisation

        - Fundamentals, stochastic gradient descent, duality and minmax opt.
        - Adaptive algorithms, interior point Method
    - PDE-constrained optimisation
        - Examples from seismic and/or acoustic imaging
        - The adjoint method
        - Regularisation
    - Non-convex optimisation
        - Motivation and fundamentals: saddle points, local minima, Hessian descent, etc.
        - Global v/s non-convex settings: Bayesian optimisation
            - Case studies in inverse problems
        - Overparameterization, Optimisation in (non-convex) high-dimensional spaces, case study: deep learning
    - Applications: optimisation in science and engineering

Literature:

- Å. Björck, Numerical Methods in Matrix Computations, Texts in Appl Maths, Vol 59, 2015
- Y. Saad. Iterative Methods for Sparse Linear Systems, 2nd ed. SIAM, 2003
- G. Golub and Ch. Van Loan. Matrix Computations, 4th ed., 2013
- J. Norcedal and S. Wright, Numerical Optimization, Springer, 1999
- Y. Nesterov, Lectures on Convex Optimization, Springer, 2018
- A. FIchtner, Full Seismic Waveform Modelling and Inversion, Springer, 2011

Contact persons:

- Roman Iakymchuk (roman.iakymchuk@it.uu.se)
- Prashant Singh (prashant.singh@scilifelab.uu.se)
- Martin Almquist (martin.almquist@it.uu.se)